

Science Goal Monitor - science goal driven automation for NASA missions

Anuradha Koratkar^a, Sandy Grosvenor^b, John Jung^c, Melissa Pell^b, David Matusow^c, Charles Bailyn^d

^aGoddard Earth Science and Technology Center, University of Maryland Baltimore County,
3.002 South Campus 1000 Hilltop Circle, Baltimore MD 21250

^bScience Systems and Applications, Inc., 10210 Greenbelt Road, Suite 500, Lanham, MD 20706

^cAdvanced Architectures and Automation Branch, NASA/Goddard Space Flight Center,
Building 23, Code 588, Greenbelt MD 20771

^dDepartment of Astronomy, Yale University, P.O. Box 208101, New Haven CT 06520-8101

ABSTRACT

Infusion of automation technologies into NASA's future missions will be essential because of the need to: (1) effectively handle an exponentially increasing volume of scientific data, (2) successfully meet dynamic, opportunistic scientific goals and objectives, and (3) substantially reduce mission operations staff and costs. While much effort has gone into automating routine spacecraft operations to reduce human workload and hence costs, applying *intelligent automation to the science side*, i.e., science data acquisition, data analysis and reactions to that data analysis in a timely and still scientifically valid manner, has been *relatively under-emphasized*.

In order to introduce science driven automation in missions, we must be able to: capture and interpret the science goals of observing programs, represent those goals in machine interpretable language; and allow spacecrafts' onboard systems to autonomously react to the scientist's goals. In short, *we must teach our platforms to dynamically understand, recognize, and react to the scientists' goals*.

The Science Goal Monitor (SGM) project at NASA Goddard Space Flight Center is a prototype software tool being developed to determine the best strategies for implementing science goal driven automation in missions. *The tools being developed in SGM improve the ability to monitor and react to the changing status of scientific events*. The SGM system enables scientists to specify what to look for and how to react in descriptive rather than technical terms. The system monitors streams of science data to identify occurrences of key events previously specified by the scientist. When an event occurs, the system autonomously coordinates the execution of the scientist's desired reactions. Through SGM, we will improve our understanding about the capabilities needed onboard for success, develop metrics to understand the potential increase in science returns, and develop an "operational" prototype so that the perceived risks associated with increased use of automation can be reduced.

SGM is currently focused on two collaborations:

1. Yale University's [SMARTS](#) (Small and Moderate Aperture Research Telescope System) observing program - Modeling and testing ways in which SGM can be used to improve scientific returns on observing programs involving intrinsically variable astronomical targets.
2. The [EO-1](#) (Earth Observing - 1) mission - Modeling and testing ways in which SGM can be used to autonomously coordinate multiple platforms based on a set of scientific criteria.

In this paper, we will discuss the status of the SGM project focusing primarily on our progress with the SMARTS collaboration.

Keywords: science optimization, data communications, onboard analysis, spacecraft autonomy, science goals

1. INTRODUCTION

NASA science missions have traditionally operated on the assumption that we can only manage scheduling priorities and scientific processing on the ground with significant human interaction, and that all scientific data must be

downloaded and archived regardless of its scientific value. However, increases in onboard processing and storage capabilities of spacecraft, as well as increases in rates of data accumulation will soon force NASA operations staff and scientists to re-evaluate the assumption that all science must be done on the ground. In order to take advantage of these new in-flight capabilities, improve science return and contain costs, we must develop strategies that will help reduce the perceived risk associated with increased use of automation in all aspects of spacecraft operations.

An important aspect of science operations is the ability to respond to science driven events in a timely manner. For such investigations, *we must teach our observing platforms to intelligently achieve the scientists' goals*. The first steps in this direction have just started. Ground based robotic telescopes have started scanning the skies, and in space the [SWIFT](#) spacecraft will have the capability of detecting gamma-ray bursts (GRBs) and then executing a rapid autonomous slew to quickly re-point the spacecraft and bring the burst within the narrow field-of-view of the X-ray and UV/Optical telescopes on the same platform.

The Science Goal Monitor ([SGM](#); <http://aaa.gsfc.nasa.gov/SGM>) is a prototype software tool being developed to explore strategies for implementing science goal driven operations for multiple sensors/platforms. Our project has both space and earth science collaborations. In the space science domain, we are implementing a prototype for dynamic automated reactions to intrinsically varying astronomical phenomenon using one of the Small and Moderate Aperture Research Telescope System (SMARTS) telescopes. Our primary objectives with the SGM/SMARTS collaboration is to prototype a small set of test scenarios in an environment that is scientifically and operationally realistic, and one in which we can measure the effectiveness of SGM. In the Earth science domain we are collaborating with the Earth Observing – 1 (EO-1) team to evaluate how multiple sensors can react dynamically to obtain rapid observations of evolving earth science events. *By developing and testing a prototype in an operational environment, we are in the process of establishing metrics to gauge the success of automating science campaigns*. In the following sections we discuss the status of each of these collaborations and the lessons learned so far, focusing primarily on the SMARTS astronomical collaboration.

2. SCIENCE GOAL MONITOR (SGM)

SGM is a set of tools that will have the ability to capture the underlying science goals of an observation, translate them into a machine interpretable format, and then autonomously recognize and react in a timely fashion when goals are met. SGM provides users with visual tools to capture their scientific goals in terms of measurable objectives. The scientists specify what to look for and how to react in descriptive rather than technical terms. The SGM system then autonomously monitors the data stream in near-real time to see if these goals (events previously specified by the scientist) are being met. When an event occurs, the system autonomously coordinates the execution of the scientist's desired reactions between different observatories or satellites. Our prototype is designed for use in a distributed environment where some analysis might be performed onboard a spacecraft, while other analyses might be performed on the ground. At the previous conference on Optimizing Scientific Return for Astronomy, we presented our rationale for the SGM project and the preliminary architecture [1].

An important research goal of SGM is to convert science goals into machine interpretable format. Given the current state of natural language technology, this is a very complex problem. However, for time-variable phenomenon, goals can be defined as measurable objectives with contingency plans for follow-on work. Hence, SGM is designed to be adaptable to many different science domains that require rapid response to fast temporal events such as gamma ray bursts or hazardous events such as forest fires, floods and volcanic eruptions.

2.1. Space Science Collaboration

To determine the effectiveness of SGM in the astronomy domain, SGM is collaborating with SMARTS, to prototype and test dynamic scheduling capabilities. Although SMARTS is a ground-based consortium, this collaboration will allow us to better understand and measure the risks and rewards of dynamic scheduling and improve the likelihood of successfully using it on space-based missions.

The SMARTS program was in its first observing season during the spring of 2003. This has been an advantage for SGM in that the SMARTS team has not had a well-established suite of operational processes that are difficult to change. The disadvantage has been that the SMARTS team members have not had a baseline set of science programs and pre-existing statistics with which SGM can be compared. Also, since the paradigm of defining an observing program, as a

set of science goals and reactions is unusual from the scientist's perspective, we have opted to focus on a subset of the current SMARTS programs where the SGM paradigm is most applicable and its impact most easily measurable.

We identified three test cases that are sufficiently different and together will represent the entire range of SMARTS proposals. These cases are programmatically challenging, and will effectively test the various aspects of SGM.

- The Gamma-ray (GRB) scenario will test how SGM and the planning and scheduling system interact with each other to assist with rapid response to a GRB alert. In this scenario SGM's task will be to act as a task manager for not only the observing night, but also ultimately the entire observing season. We discuss the details of the scenario in section 2.1.1.
- The X-ray binary scenario is an ideal test case for SGM because it tests many modules in SGM. In this test scenario, SGM will monitor and track target brightness of multiple X-ray binaries to detect and capture an event/outburst. When the start of an outburst is detected, SGM will implement a number of activities. It will adjust the observing priorities of all targets in the X-ray binary program in order to best implement the revised strategy for observing the active binary. SGM will also implement the scientifically driven contingency plans so that the outburst is captured as effectively and efficiently as possible. Thus, in this scenario we will test SGM's ability to autonomously monitor data, detect an event and implement a scientist's observing goals when that event is detected.
- The Supernova scenario will test how well SGM can implement and adapt an observing strategy over a longer time period with a strict criteria. In this scenario, SGM will monitor internet resources of supernova discoveries such as Supernova Factory, LOTOSS, and amateur discoveries. When a new supernova has been detected, SGM will perform initial imaging of the supernova, and fire off high-priority data analysis on the new images. The results of the data analysis will affect the priority and observing strategies for the target over a multi-week period and depend on the nightly observing conditions in the first several nights of observing the supernova. This scenario is valuable because the observations will be constrained by night conditions and there are very strict demands on the observing strategy and associated data processing. Thus, this scenario will test SGM's ability to implement an exacting observing strategy and its ability to interact with a data analysis system for associated data processing.

2.1.1. Prototype Status

The SMARTS telescope operations process is that an astronomer at Yale University generates the observing plan or nightly schedule. The schedule is sent to Chile where telescope operators execute it. The nightly schedule remains static throughout the night because; (1) it is generated in Yale with tools that rely on manual input, and (2) there is typically not an astronomer on duty to make decisions regarding the science impacts of the schedule changes. Consequently, if there is any disruption (clouds, instrument failure, target of opportunity) the operators in Chile have almost no ability to react in a fashion sensitive to scientific priorities. This operations strategy was not amenable for SGM interaction. To be able to rapidly respond to evolving temporal events, SGM needs to operate in a live and dynamic environment. This implied that the some aspects of the SMARTS operations strategy needed to be changed. Thus, the first phase of development for the SGM and SMARTS collaborative teams have been to develop tools to allow SGM to communicate with the observing queue.

Our initial tools include two closely related modules, the *ScheduleAssistant* and *OperationsAssistant*. They provide a web-accessible, database-driven interface between the scheduling astronomer, operations teams, and SGM. The scheduling astronomer at Yale uses the *ScheduleAssistant* to develop a nightly observing schedule with priorities attached to the scheduled observations. The schedule is published to the database, and then the operations team in Chile can access the schedule using the *OperationsAssistant*. The operations staff can update the data with information on the actual observations. These simple assistants automate much of the labor-intensive "cut and paste" processes that the SMARTS team had used to develop their night schedule. The tool has been quick and simple to develop, but also has enough details that it is useful to both SGM and SMARTS. It provides an essential interface for SGM, while giving the SMARTS team a significant and immediate labor saving efficiency boost.

We are now in the process of modeling our initial scenarios and have just started to dynamically schedule the nightly schedule. This required us to replace the SMARTS scheduling scripts with a more dynamic planning and scheduling system. We have chosen NASA's [ASPEN](#) scheduling system [2], because of the modularity of its *design and ease of*

- Expanding the GRB scenario sufficiently to support the growth in GRB alerts expected when the SWIFT satellite is launched (currently planned for September, 2004).
- Increasing the capabilities of the re-planning system to allow SGM to handle the remaining scenarios discussed above. Implementing and testing these scenarios.
- Developing and analyzing metrics to evaluate the effectiveness and reliability of SGM.
- Releasing the *ScheduleAssistant* and the *OperatorAssistant* as open-source tools to the astronomical community.

The early feedback on both the *ScheduleAssistant* and the *OperatorAssistant* tools from the SMARTS team has been very favorable and we've been encouraged to consider adapting the tools for use by other small and medium-size observing teams. Since the SGM team lacks the resources to extend the tools beyond the immediate needs of the SGM/SMARTS collaboration, we hope that open-sourcing the tool will allow other observatories to use it and adapt it and let it further evolve.

2.2. Earth Science Collaboration

Using NASA's Earth Observing-1 (EO-1) satellite (<http://eo1.gsfc.nasa.gov>) and Earth Observing Systems' *Aqua/Terra* spacecrafts' MODIS instrument (<http://modis.gsfc.nasa.gov>), we have recently completed a series of prototype demonstration tests of varied phenomenon such as: forest fires, volcanoes, floods and lake freezing. In these demonstrations SGM analyzed data received from the MODIS instruments with a dynamic autonomous request for higher-resolution images from the EO-1 satellite based on a set of scientific criteria. Figure 2 shows the data flow for the EO-1 forest fire demonstration. Our demonstrations so far have been relatively simple to develop and demonstrate the basic capabilities of SGM. They show the promise of coordinating data from different sources, analyzing the data for a scientifically relevant event, and autonomously updating and rapidly obtaining a follow-on scientifically relevant image. Details of the demonstrations are discussed in [3]. *The demonstrations showed that relevant science events can be accurately detected and that the details of the follow-on observations can be correctly captured and transmitted to the appropriate location autonomously. Further, the demonstrations show that existing assets, which were not built for autonomy, or flexible scheduling, and have limited onboard processing, can still be autonomously coordinated and*

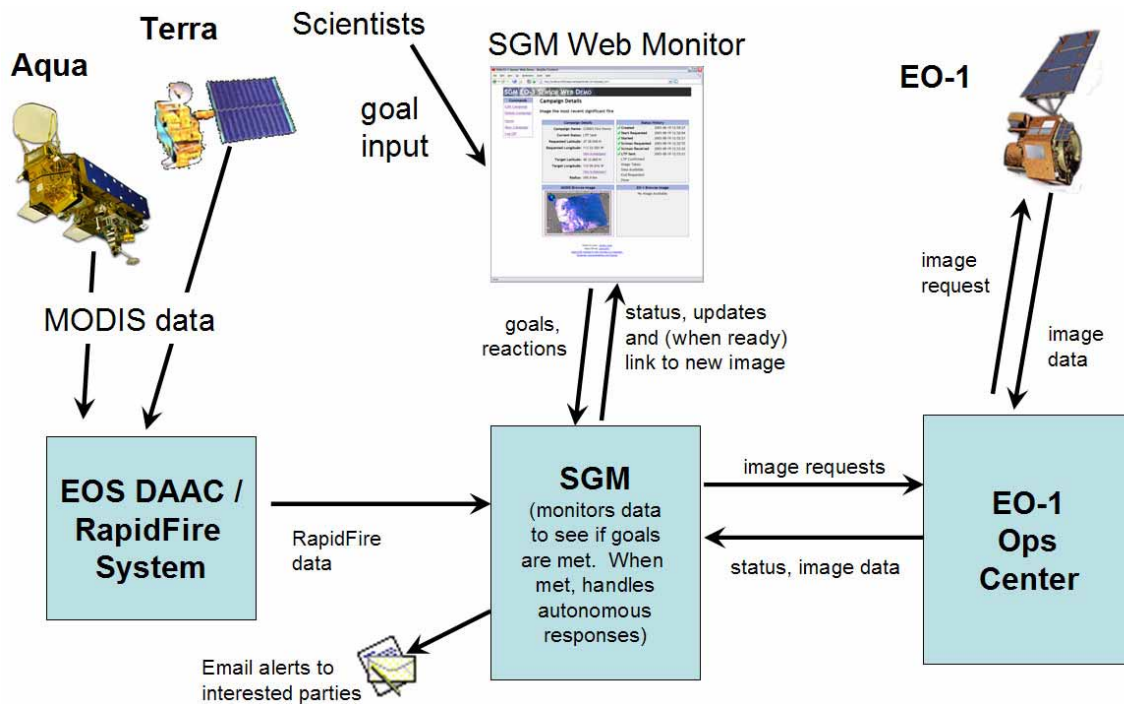


Figure 2: SGM data flow for the EO-1 forest fire demonstration

dynamically scheduled. Automation of a legacy system is easier to achieve if that automation is layered on top of the existing system. However, the limit of such automation is directly related to the flexibility of the existing system design. We are now in the process of developing new science scenarios with more complex reasoning. For example, SGM will be monitoring a prioritized list of volcanoes for new eruptions using both in-situ and satellite data. When an eruption or inflation event is detected, SGM will coordinate with EO-1 planning to automatically request a high-resolution image of the volcano area. Further, by accessing real-time weather data from the GOES satellite, SGM will coordinate “near-real time” with EO-1 planning to obtain cloud free images to maximize the scientific value of the image obtained.

3. ESTABLISHING METRICS FOR THE SGM/SMARTS COLLABORATION

Before implementing science driven automation in missions, we need to not only understand the capabilities that are needed onboard, but also realistically understand the potential increase in science returns. Therefore, developing metrics and evaluating SGM is an important activity of the SGM team. The specific productivity impacts of SGM/SMARTS that we hope to measure are:

- Reduction in the time taken to change/implement observing strategies in response to scientific events.
- Decrease in the amount of time science team members spend on maintaining science programs.
- Decrease in the response time to a scientific event.

Observatories have defined metrics in many different ways to track their various processes and contributions to the scientific community. For operational metrics, observatories often rely on statistics such as percentage of time spent exposing. On the science side, traditional metrics frequently include long-term measures such as the number of citations in refereed scientific journals. We are striving to compare effectiveness of operations with and without a prototype tool that changes observing strategy in response to science driven needs. Unfortunately we feel that many previously used metrics [4-7] would not successfully measure the effectiveness of SGM. We have therefore established a baseline set of new metrics that we hope will help quantify and demonstrate where and how SGM is most effective in improving quality and efficiency in scientific terms.

There are two types of metrics that we will use. The first are operational metrics that will quantitatively measure the amount of time spent on various tasks. If SGM is successful, the total hours spent in rote work, which deals with the planning and maintaining of a science program should reduce. The second metric is scientific success. This is a more subjective measure, yet it will be effective in measuring SGM’s successes as perceived by a scientist. This metric is important because scientists, especially astronomers, have to be convinced that intelligent automation of scientific tasks is possible; that unique science goals can be effectively and accurately captured and executed.

3.1. Operational metrics

- To track the time spent on each program by the astronomer and the observatory staff we will measure the hours spent on an observing program by type of work being performed. This will include tasks such as:
 - o Observation planning (establishing, then maintaining a science program's detailed observation plan)
 - o Preliminary data processing of observations to determine if there is an event
 - o Scientific analysis of observations
 - o Other administrative tasks not mentioned above
- To track the effectiveness of SGM in responding to events we will determine the following measures.
 - o Detection time, which we define as the time interval between acquiring an observation containing a science event and the detection of that event.
 - o Reaction time, which we define as the time interval between detection of an event and the acquisition of the first image of a revised observing strategy.
- To track improvements in observatory operations and to determine if SGM has maximized science returns we will track the
 - o Number of high priority observations successfully re-scheduled and obtained after a disruption of the schedule compared to the number of high priority observations in the original schedule.

- o Ratio of originally requested observations to the number of observations actually scheduled for each program.
- o Ratio of the original number of scheduled observations to the number of actual observations obtained. We will pay attention to the fact that as a program is executed, the campaign strategy and number of observations may change.

3.2. Scientific metrics

Since these metrics are very subjective, we felt that the SGM user (PI's for our three prototype test scenarios) should define a measure of success that is unique to their type of program. SGM will be considered to be successful if it meets the user's predetermined base value for the user-defined metric. We are in the process of finalizing these metrics.

Baseline measurements of the operational metrics began with the fall 2003 semester of SMARTS observing prior to introducing SGM into the process. Since, our first priority has been to get SGM to communicate with the observing queue in a live, dynamic basis we have yet to start collecting metrics to evaluate SGM.

4. LESSONS LEARNED

4.1. Flexible user interface vs. templates

When this research began, our intent was to develop a flexible user interface that allowed the scientist to specify a wide range of science goals. We developed a prototype that used visual programming concepts to provide a set of graphical building blocks with which the user could construct goals. While it was very flexible, the overwhelming reaction was that it was much too difficult to construct science goals. We discovered that the majority of science goals could be represented by a set of adjustable templates. This led to the abandonment of this first prototype and caused us to rethink our approach. Instead of a generic infinitely flexible user interface, we now present the user with a list of observation templates that the user can customize. Each template defines a typical kind of observation that the user might want to perform. These templates include parameters that the user can specify to customize their campaign, while the main template structure is fixed. While this approach does not allow the scientist to express all possible science goals, it does satisfy the majority of them, and since the system is quite easy to use, reaction to this new approach has been very positive.

The underlying architecture of the template system is based on a combination of XML, Java servlets, and the Java Struts interface. Once we define a template of an observing campaign in XML, the servlets interface can read the template and automatically generate a web-based form that an end-user can fill out to establish a specific campaign instance. The user can request that SGM begin executing the campaign and monitor the status of the campaign via the same standard web interface.

4.2. Automation in parts

A fully automated/autonomous system is very difficult to build because the environment in which this system must operate is constantly changing. This change comes about because the nature of the human interaction with the system changes. Thus, an autonomous system not only must have the ability to adapt to changes, but must also have a well-defined plan by which a human can interact with the system. Therefore, in the SGM project, we have not tried to automate the entire science operations process, but have taken the approach that human interaction is essential to the autonomous system. We are dividing up the problem of flexible scheduling and autonomously reacting to science driven events into a number of smaller more realistically achievable goals.

A significant challenge has been in deciding which processes should be automated, and which should remain manual. Certainly those that require exact mathematical computation can and should be automated, but when we considered aspects that involved human decision-making, there have been concerns among the scientists. For example, in assigning priorities to observations for the purpose of optimizing the total observing run, a scientist will consider many factors, and essentially determine the scientific worth of an observation, which is fundamentally a subjective task. We decided to attempt to automate as much as possible of this process, but insert the human in the automation loop so that the human can provide their own subjective input into the process where necessary. For example, in the Earth science

demonstrations we have automated the mundane tasks of (1) monitoring data, (2) generating the technical details for the follow-on observation, (3) requesting the follow-on observations, and (4) tracking the status of the observation. The scientists are free to focus on the higher-end subjective science analysis for which they have studied and trained. An additional benefit of this strategy was that it alleviated the concerns of the users. The purpose of the system was not to automate and thus replace their jobs, but instead to automate only the tedious aspect of their jobs, freeing them to concentrate on the more interesting aspects.

The next problem is the task of developing a dynamic schedule or observing plan. Rapidly or semi-autonomously responding to scientific events can be very disruptive to a spacecrafts' observing efficiency. In the SMARTS collaboration we are focusing on this aspect of the project. The observing schedule is currently developed manually on a daily basis by juggling various requirements and priorities of the different observing programs. Once the nightly observing begins, the schedule remains static. If a disruption occurs (such as cloud cover, instrument failure etc.), once observing resumes the operators simply skip observations whose scheduled start time has passed. In our first phase of enabling autonomous scheduling, we are simply aiming to improve the science returns for re-scheduling the remainder of a single night after a disruption. Dynamically rescheduling the observing plan over a short timescale will be very useful in the context of onboard scheduling. This is a far simpler challenge to model, execute, and compare against the static schedule.

4.3. Knowledge capture

Once we established that we would structure the system using a menu of science templates, we have sought to build those templates in cooperation with our scientist users. This process primarily consists of a series of face-to-face interviews with the scientists, followed by emails to clarify particular issues. The first task is to identifying the set of science phenomena that we would support in our prototype. For each type of phenomenon, we develop a story for what the scientists want to observe and how they would go about observing it. Each story contained branches where different things would occur if certain criteria were met in the science data. One challenge was coming to an agreement on the story, particularly because the different scientists that we interviewed did not always agree. It is a very iterative process where we collect each scientist's version of the story, identify differences, and then clarify those differences with each scientist. In some cases, we realized the need for a new user adjustable parameter where each scientist might tweak a value in the story. Fortunately, this has not been a great problem for the demonstrations in Earth science, because the differences between the scientists' stories have generally been relatively minor. We are just starting this process for our SMARTS collaboration.

5. CONCLUSIONS

Developing a spacecraft with flexible scheduling, which can autonomously react to science driven events is not just a leap forward in automation, but a large change in operations paradigm. There are a number of challenges that need to be overcome to change the present NASA mission operation strategies. The SGM project is a proof-of-concept effort to address these challenges.

In SGM we are developing an interactive distributed system that will use onboard processing and storage combined with event-driven interfaces with ground-based processing and operations, to enable fast reaction to time-variable phenomena. We are currently developing prototypes and evaluating the effectiveness of the system. Although we have not completed our project, we have had some success.

- The SMARTS team got a significant operational efficiency boost from the two simple tools that we developed for SGM integration into their system. The team is looking forward to implementing the test scenarios and evaluating the effectiveness of SGM.
- Our dynamic science analysis and autonomous multi-sensor coordination has been highly appreciated in the Earth science domain and scientists are looking forward to attempting more complex problems. SGM demonstrated a multi-sensor coordination and a more responsive science operations paradigm with dramatic gains in science returns. For example, in our forest fire demonstration, SGM coordination and analysis provided new data to the US Forestry Service within 48 hours, compared to a typical lead-time of up to 14 days for preplanned observations.

ACKNOWLEDGEMENTS

NASA Code R funds this work under the Computing, Information and Communication Technologies (CICT) program.

REFERENCES

- [1] Koratkar, A., Grosvenor, S., Jones, J. E., Memarsadeghi, A., & Wolf, K. R., *Science Goal Driven Observing: A Step towards Maximizing Science Returns and Spacecraft Autonomy*, 2002, SPIE 4844, 250.
- [2] Chien, S., Knight, R., Stechert, S., Sherwood, R., & Rabideau, G., *Integrated Planning and Execution for Autonomous Spacecraft*, 1999, Proceedings of the IEEE Aerospace Conference (IAC), Aspen, CO.
- [3] Koratkar, A., Jones, J. E., Jung, J., & Grosvenor, S., *Science Goal Driven Automation for NASA Missions: The Science Goal Monitor*, 2004, AAAI paper.
- [4] Doxsey, R.E., Downes, R., Lallo, M., & Reinhart, M., *Metrics for Hubble Space Telescope science operations*, 2002, SPIE 4844, 104.
- [5] Adler, D. S., Taylor, D. K., & Patterson, A. P., *Twelve years of planning and scheduling the Hubble Space Telescope: process improvements and the related observing efficiency gains*, 2002, SPIE 4844, 111.
- [6] Benn, C.R., & Talbot, G., *Telescope performance metrics*, 2002, SPIE 4844, 122.
- [7] Neilsen, E., & Boroski, W. N., *Sloan Digital Sky Survey observing time tracking and efficiency measurement*, 2002, SPIE 4844, 130.